

What Is Claimed Is:

1. A processor-implemented method of routing network traffic through fault zones, comprising:
 - 5 identifying a path from a first network node to a second network node;
identifying a set of fault zones through which the identified path passes;
for each fault zone in the set of fault zones, assigning as a zone weight the number of paths from the first network node to the second network node that include said fault zone;
 - 10 calculating a path weight for the identified path, wherein said path weight is equal to the sums of said zone weights for each fault zone included in the identified path; and
selecting the identified path as the current path for network traffic from the first node to the second node.
- 15 2. The method of claim 1, further comprising:
identifying a new path from the first network node to the second network node;
assigning zone weights to each fault zone in the new path;
20 calculating a new path weight for the new path; and
if the new path weight is lower than said path weight for the identified path, selecting the new path as the current path for network traffic from the first node to the second node.
- 25 3. The method of claim 1, wherein:
the first network node is identified by a first identifier;
the second network node is identified by multiple identifiers, including a

second identifier;

selecting the identified path as the current path for network traffic from the first node to the second node comprises selecting the identified path the current path for network traffic from the first identifier to the second identifier; and

5 paths other than the identified path are selected as the current paths for network traffic from the first identifier to the multiple identifiers other than the second identifier.

4. A computer readable medium storing instructions that, when
10 executed by a computer, cause the computer to perform a method of routing network traffic through fault zones, the method comprising:

identifying a path from a first network node to a second network node;

identifying a set of fault zones through which the identified path leads;

for each fault zone in the set of fault zones, assigning as a zone weight the
15 number of paths from the first network node to the second network node that include said fault zone;

calculating a path weight for the identified path, wherein said path weight is equal to the sums of said zone weights for each fault zone included in the identified path; and

20 selecting the identified path as the current path for network traffic from the first node to the second node.

5. A processor-implemented method of determining routing between nodes in a subnet, comprising:

25 identifying multiple fault zones in the subnet, each fault zone comprising one or more components of the subnet;

configuring a central subnet manager to manage routing between nodes in

the subnet;

identifying a set of paths from a first node having a first identifier to a second node having multiple identifiers, including a second identifier, wherein traffic is deliverable to the second node using any of the multiple identifiers;

5 for each fault zone traversed by one or more of the paths, establishing a zone weight based on the number of paths from the first node to the second node that traverse said fault zone;

for each path in the set of paths, establishing a path weight from the sums of the zone weights for each fault zone traversed by said path; and

10 for each of the multiple identifiers of the second node, selecting as the current path from the first identifier to said identifier, from said set of paths, the path having the best path weight.

6. A computer readable medium storing instructions that, when
15 executed by a computer, cause the computer to perform a method of determining routing between nodes in a subnet, the method comprising:

identifying multiple fault zones in the subnet, each fault zone comprising one or more components of the subnet;

20 configuring a central subnet manager to manage routing between nodes in the subnet;

identifying a set of paths from a first node having a first identifier to a second node having multiple identifiers, including a second identifier, wherein traffic is deliverable to the second node using any of the multiple identifiers;

25 for each fault zone traversed by one or more of the paths, establishing a zone weight based on the number of paths from the first node to the second node that traverse said fault zone;

for each path in the set of paths, establishing a path weight from the sums

of the zone weights for each fault zone traversed by said path; and

for each of the multiple identifiers of the second node, selecting as the current path from the first identifier to said identifier, from said set of paths, the path having the best path weight.

5

7. A processor-implemented method of routing traffic in a subnet based on mean time between failure (MTBF) data of components within the network, the method comprising:

for each of multiple paths from a first node identifier to a second node
10 identifier, wherein the first node identifier identifies a first subnet node and the second node identifier identifies a second subnet node:

for each of multiple subnet components included in said path,
identifying MTBF data of said component; and

aggregating the MTBF data for the multiple subnet components to
15 calculate a path weight for said path; and

based on the path weights, selecting one of the multiple paths as the current path for subnet traffic from the first node identifier to the second node identifier.

20 8. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of routing traffic in a subnet based on mean time between failure (MTBF) data of components within the network, the method comprising:

for each of multiple paths from a first node identifier to a second node
25 identifier, wherein the first node identifier identifies a first subnet node and the second node identifier identifies a second subnet node:

for each of multiple subnet components included in said path,

identifying MTBF data of said component; and
aggregating the MTBF data for the multiple subnet components to
calculate a path weight for said path; and
based on the path weights, selecting one of the multiple paths as the
5 current path for subnet traffic from the first node identifier to the second node
identifier.

9. A processor-implemented method of routing traffic in a subnet
based on one or more selected characteristics, the method comprising:
10 for each of multiple paths from a first node identifier to a second node
identifier, wherein the first node identifier identifies a first subnet node and the
second node identifier identifies a second subnet node:
for each of multiple subnet components included in said path,
identifying values for one or more characteristics of said component; and
15 aggregating the values for said one or more characteristics for the
multiple subnet components to calculate a path weight for said path; and
based on the path weights, selecting one of the multiple paths as the
current path for subnet traffic from the first node identifier to the second node
identifier.

20
10. The method of claim 9, wherein:
each of the multiple subnet components comprises a communication link;
and
a first characteristic of a subnet component comprises a link speed of the
25 communication link.

11. The method of claim 9, wherein a first characteristic of a subnet

component comprises a hop count between the first subnet node and the subnet component.

12. The method of claim 9, wherein a first characteristic of a subnet
5 component comprises a cost associated with the subnet component.

13. The method of claim 9, wherein a first characteristic of a subnet component comprises a quality of service offered by the subnet component.

10 14. A system for determining routing through a subnet comprising multiple fault zones, comprising:

a network node configured to interface a client computing device with the subnet, wherein each node is identifiable by one or more node identifiers; and

a client computing device, comprising:

15 a processor; and

a subnet manager module configured to determine routing between a first node and a second node in the subnet, wherein the first node is addressable by a first identifier and the second node is addressable by multiple identifiers, including a second identifier;

20 wherein said subnet manager determines routing between the first node and second node by:

for each fault zone in the subnet traversed by a path from the first node to the second node, calculating a zone weight based on the number of paths from the first node to the second node that traverse said fault zone;

25 for each of the paths from the first node to the second node, calculating a path weight based on the sums of said zone weights for the fault zones traversed by said path; and

selecting as the current path from the first identifier to the second node identifier the path from the first identifier to the second identifier having the best path weight.

5 15. The system of claim 14, wherein the client computing device further comprises:

 a memory configured to store path weights of current paths between multiple pairs of node identifiers.

10 16. The system of claim 15, wherein said memory is further configured to store, in association with each of the current paths, zone weights for fault zones traversed by the current path.

 17. The system of claim 14, wherein said subnet manager is further
15 configured to disseminate routing information to a plurality of nodes in the subnet, said routing information including said current path from the first identifier to the second identifier.

 18. A system for determining routing through a subnet based on one or
20 more characteristics of subnet components, the apparatus comprising:

 a network node configured to interface a client device with the subnet; and
 a client computing device, comprising:

 a processor; and

 a subnet manager module configured to determine routing between
25 a first node and a second node in the subnet by:

 for each of multiple subnet components traversed by a path
 from the first node to the second node, identifying values of one or

- more characteristics of said subnet component;
- for each path from the first node to the second node,
establishing a path weight corresponding to the combined values of
the one or more characteristics for subnet components traversed by
said path; and
- selecting as the current path from a first identifier
associated with the first node to a second identifier associated with
the second node, the path having the best path weight.
19. The system of claim 18, wherein the subnet components comprise
one or more of:
- a network node; and
a link between two network nodes.
20. The system of claim 18, wherein a first characteristic of a subnet
component comprises a mean time between failures (MTBF) of the subnet
component.
21. The system of claim 18, wherein:
- each of the multiple subnet components comprises a communication link;
and
a first characteristic of a subnet component comprises a link speed of the
communication link.
22. The system of claim 18, wherein a first characteristic of a subnet
component comprises a hop count between the first node and the subnet
component.

23. The system of claim 18, wherein a first characteristic of a subnet component comprises a cost associated with the subnet component.

5 24. The system of claim 18, wherein a first characteristic of a subnet component comprises a quality of service offered by the subnet component.